

Output Market Segmentation and Productivity Heterogeneity

Chad Syverson
University of Maryland
syverson@econ.umd.edu

Abstract

Recent empirical investigations have shown enormous plant-level productivity heterogeneity, even within narrowly defined industries. Most of the theoretical explanations for this have focused on factors that influence the production process, such as idiosyncratic technology shocks or input price differences. I claim that characteristics of the output demand markets can also have predictable influences on the plant-level productivity distribution within an industry. Specifically, an industry's degree of output market segmentation (i.e., the substitutability of one plant's output for another's in that industry) should impact the dispersion and central tendency of the industry's plant-level productivity distribution. I test this notion empirically by seeing if measurable cross-sectional variation in market segmentation affects moments of industry's plant-level productivity distribution moments. I find significant and robust evidence consistent with this notion.

I am grateful to John Haltiwanger and John Shea for their comments, and I thank Glenn Ellison for providing the industry geographic concentration data used in the paper. The research in this paper was conducted while the author was a research associate at the Center for Economic Studies, U.S. Bureau of the Census. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census or the Center for Economic Studies.

Output Market Segmentation and Productivity Heterogeneity

Recent empirical explorations into plant-level productivity levels and growth rates have consistently shown substantial heterogeneity across plants, even within what often seem to be narrowly defined (e.g., four-digit SIC) industries. A host of theoretical work has arisen in an attempt to explain the sources of such diversity. The core explanation advanced by these theories is that plants are subject to idiosyncratic shocks to their production fundamentals. For example, across-plant shocks to wages or capital costs can lead to plant input-purchase decisions that vary significantly even under an identical production technology. Another (not mutually exclusive) possibility is that idiosyncratic technology/efficiency shocks drive output and productivity variations—through entry and exit processes, for instance. Furthermore, while these shock processes lead directly to broad across-plant differences, they can also create asymmetric allocations of productive resources that create additionally divergent outcomes among establishments. Disparate investment patterns (caused by responses to idiosyncratic shock processes) lead to a distribution of capital vintages that in turn yields heterogeneous production patterns. Differences in the mix of worker skills—possibly related to capital specificity—also affect the distribution of inputs and outputs across plants.

A useful question to consider when studying within-industry productivity heterogeneity is how such wide dispersions can exist in equilibrium. We would expect a long-run tendency for output to be reallocated to more productive plants, as they can produce output at lower-cost than industry rivals, and hence grab additional market share by undercutting their opponents' prices without sacrificing profits. Indeed, if such a process were to continue unabated, industry equilibrium would expectedly be characterized by a degenerate plant-level productivity distribution within the industry; all operating plants would have the same productivity level. With a globally increasing returns technology, the industry output distribution would be degenerate as well—one plant would produce the entire industry output. Under technologies exhibiting constant returns or standard U-shaped average costs, industry output would be symmetrically distributed among several equally productive plants operating at the maximum efficient scale. Of course, such output and productivity distributions are not observed in the data; as mentioned above, the overwhelming weight of empirical evidence indicates widely varying productivity

levels across plants within nearly every industry. I have already enumerated several possible explanations for this. The great weight of these, and of the theoretical heterogeneity literature in general, focus on the influence of supply-side (production) factors. In this chapter I turn my attention to a traditionally overlooked source of persistent heterogeneity: demand-side (output market) forces. I focus specifically on the role that product differentiation and segmented output markets play in determining the equilibrium distribution of plant-level productivity within an industry.

The intuition behind the influence of output market segmentation on an industry's productivity distribution is simple. If purchasers could perfectly substitute the output of one plant for another's within a given industry, all demand would go to the producer charging the lowest price. Producers engaging in Bertrand-type price competition would be narrowed down to the plant (or plants, if decreasing returns become a factor) that can produce at the lowest cost. Abstracting from input market differences across plants, this will be the establishment requiring the least inputs for a given amount of output: the highest productivity plant. We do not see such "knife-edge" scenarios in practice, of course, because the output of one industry establishment is not perfectly substitutable for the output of another. Several factors that contribute to output market segmentation interact to prevent such an outcome. Such departures from perfect output homogeneity within an industry allow less productive plants to survive even in long-run equilibrium. This results in a non-degenerate plant-level productivity distribution in the industry.

The sources of output market segmentation are manifold. One origin is physical product differentiation. There can be considerable variety in the physical characteristics of output even within a narrowly defined four-digit industry. The auto industry (SIC 3711) is an obvious example of a narrowly defined sector (in traditional industry taxonomies, at least) that produces output with a broad range of physical qualities. Idiosyncratic consumer preferences favoring one quality over another mean that some production plants can remain operational even if they are less physically efficient than their industry counterparts, if they produce output which appeals to certain customer segments. Staying with the auto industry example, plants which produce niche-market specialty vehicles will almost surely have higher unit costs than the plants producing vehicles for the masses. However, many such niche plants survive (and indeed thrive) because the physical qualities of their output appeal to a certain segment of

customer. Such a scenario would not be possible in a world with homogeneous, perfectly-substitutable industry output.

Segmented output markets can still exist even when industry products are physically identical. Transport costs create barriers to perfect substitutability. Consider the manufactured ice industry (SIC 2079). It is highly unlikely that the physical characteristics of output varies much from plant to plant in this industry. However, the obvious transport barriers make manufactured ice in one locale an imperfect substitute for the same product in another. So it is not surprising that we would observe a range of plant productivity levels in equilibrium, as more efficient plants would not be able to take market share from less efficient competitors if they are located sufficiently far away (implying that they should not in truth be considered competitors!).

Branding and advertising can also lead to consumers perceiving physically identical products as being less than interchangeable. Name-brand pharmaceuticals still sell despite the presence of chemically identical generic alternatives, for instance. With sufficient brand identity, a plant can continue to operate even if it is less efficient than its industry competitors.

Real or perceived differences in services bundled with products, such as delivery speed and reliability, documentation, and product support, also contribute to output market segmentation. Finally, an array of intangible factors such as specific, history-laden relationships between plants and their customers; interpersonal customer-manager interaction; and other assets of goodwill make costless substitution of another establishment's output impossible.

All of these influences, either alone or in combination, create non-degenerate distributions in measured productivity levels among plants in an industry. One would expect certain descriptive moments of the productivity distribution to vary with the magnitude of the factors affecting output substitutability in that industry. Two such implications present themselves as intuitively descriptive and testable, and as such are the subject of my investigation. The first is that the dispersion of plant productivity distribution in an industry should increase as its degree of output market segmentation grows: less substitutability among plant outputs should lead to greater variation in the productivity of operating plants in equilibrium. The second is that greater segmentation in output markets should decrease the central tendency of an industry's plant-level productivity distribution relative to other

industries. The insight here of course is that substitutability allows more productive plants to eliminate less efficient ones in equilibrium. More segmentation means that low productivity plants will still be able to produce in the long-run.

To empirically test these hypotheses, I combine a 1977 cross section of all U.S. manufacturing plants, the 1977 Census of Manufactures (CM), with an assortment of other industry-level data on measurable contributing factors to output substitutability. I test for the influence of output market segmentation across four-digit Standard Industrial Classification (SIC) industries on their productivity distribution moments of dispersion and central tendency. To preview the results, I find that influences on the substitutability between plant-level outputs do indeed affect the within-industry dispersion of plant productivity measures in the expected direction, even after controlling for several other plausible causes of heterogeneity. Furthermore, the same substitutability-changing factors also explain an economically important portion of the variance in average total factor productivity levels across industries.

I. Theoretical Motivation

To formalize the story linking output substitutability and the local productivity distribution, I require a theoretical framework that incorporates heterogeneous producers and contains some notion of substitutability. Further, it should allow the endogenous determination of the equilibrium plant productivity distribution, and offer testable implications as to the nature of the distribution as substitutability varies. A model employed by Melitz (1999) in another context meets these requirements. This model, which I briefly describe below to motivate my empirical tests, nicely incorporates these items and will serve as a theoretical foundation for my empirical work. Because the model has been thoroughly discussed in Melitz's work, I forgo much of the formal analysis here and focus on intuitive discussion.

The model frames the allocation of market demand across producers in the familiar Dixit-Stiglitz (1977) structure. This setup allows interplant output substitutability to enter explicitly into the model, and offers an analytically tractable way to incorporate heterogeneous producers into an equilibrium market structure.

A representative consumer in the modeled local market has C.E.S. preferences over a

continuum of goods (producers) indexed by i :

$$U = \left[\int_{i \in I} q(i)^r di \right]^{\frac{1}{r}} \quad (1)$$

This is the familiar Dixit-Stiglitz utility function. The extent to which the output of industry plants varies is captured by the parameter r , assumed to lie on the interval (0,1).

As is well known, the utility function above implies an elasticity of substitution between any two goods equal to

$$s = \frac{1}{1-r} \quad (2)$$

This value is always greater than one because of the assumption that different varieties are substitutes ($0 < r < 1$). It is easy to show that the quantity ratio between any two varieties is determined completely by the price ratio:

$$\frac{q(i_1)}{q(i_2)} = \left[\frac{p(i_1)}{p(i_2)} \right]^{-s} \quad (3)$$

Production requires a single input, labor, which is supplied elastically to the industry at a wage w . The production function is linear in labor and includes an overhead labor fixed cost f . Plants differ only in their marginal product of labor, embodied in the productivity value f :

$$q = f(l - f) \quad (4)$$

As Dixit and Stiglitz show, the demand structure outlined above leads to each producer pricing their output at the same markup over marginal cost. Normalizing $w = 1$, the optimal price for each producer is

$$p = \text{markup} \cdot MC = \frac{s}{s-1} \frac{w}{f} = \frac{1}{rf} \quad (5)$$

Therefore more efficient producers sell at a lower price.

Given that plant profits are equal to $\mathbf{p} = r - l = pq - l$, substitution of the production function and optimal pricing rule into this expression yields establishment profits of

$$\mathbf{p} = \frac{r}{s} - f \quad (6)$$

Note that the above production structure will yield an output and revenue dispersion across

$$\frac{q(f_i)}{q(f_j)} = \left(\frac{f_i}{f_j} \right)^s \quad \text{and} \quad \frac{r(f_i)}{r(f_j)} = \left(\frac{f_i}{f_j} \right)^{s-1} \quad 5 \quad (7)$$

establishments governed by plants' relative efficiencies:

Thus more productive plants produce more and have higher total sales. Notice, too, that as the elasticity of substitution s increases, the output and revenue distributions skew further toward high-efficiency (low-price) producers. In the limiting case of perfect substitutability, $s \rightarrow \infty$, the most efficient producer accounts for the entire output of the industry, as discussed above.

Equilibrium will be characterized by a mass M of establishments with a productivity distribution over some portion of $(0, \infty)$. Now define a market quantity index $Q = U$ and a price index P such that

$$P = \left[\int_0^\infty [p(f)]^{1-s} M m(f) df \right]^{\frac{1}{s-1}} \quad (8)$$

where $p(f)$ is the price set by producers with productivity f and $m(f)$ is the probability density function of operating plants' productivity distribution. This can be rewritten by using the optimal pricing rule as

$$P = M^{\frac{1}{1-s}} p(\tilde{f}) \quad (9)$$

$$\tilde{f} = \left[\int_0^\infty f^{s-1} m(f) df \right]^{\frac{1}{s-1}} \quad (10)$$

where

This moment of the productivity distribution, which I shall refer to as average productivity (it can be shown to be a quantity-weighted average of plant productivity levels), has convenient properties for aggregation. As Melitz (1999) shows, total market revenue R (and profits Π) can also be expressed as simple relationships between the mass of producers and the revenue (profits) of the establishment with this average productivity level. That is,

$$R = \int_0^\infty r(f) M m(f) df = M r(\tilde{f}) \quad \text{and} \quad \Pi = \int_0^\infty p(f) M m(f) df = M p(\tilde{f}) \quad (11)$$

It is easy to see from these expressions that the plant with average productivity will have revenue and profits equal to the market averages.

Entry and exit processes take place in a dynamic framework. An infinite number of potential entrants exist, and can enter at any time if they pay a sunk labor cost of entry f_e . All entrants discover their productivity value f , drawn from a common distribution $g(f)$ over $(0, \infty)$, upon paying the entry

cost. Plant productivity levels are constant over time. After discovery of their efficiency level, plants can choose to either produce—in which case they are subject to the additional fixed cost of production f —or exit immediately at no additional cost. A steady-state equilibrium with a constant productivity distribution is maintained by the assumption that any producing plant is subject to an exit-forcing shock with exogenous probability \mathbf{d} , which does not vary with the productivity level. In equilibrium, the mass of successful entrants equals the exiting mass.

Both $g(\mathbf{f})$ and \mathbf{d} are exogenous, but there is considerable flexibility allowed in their form and magnitude, improving the generality of results. Slightly more restrictive is the assumption that the probability of exit is irrespective of an establishment's productivity level. This notion seems to be contrary to the plant-level evidence of entry and exit (see Dunne, Roberts, and Samuelson (1989), for instance). However, if one interprets the initial (endogenous) entrant's produce/exit choice as occurring over a short period of production in reality, rather than immediately as in the model, it can explain the tendency seen in empirical work for younger and less efficient plants to be more likely to exit.

While the productivity distribution prior to entry is exogenous, the distribution of actively operating plants is determined within the model. This is so because, as can be shown, there is always a positive productivity level where plant profits are zero. And because plant productivity levels are constant, no producer less efficient than this cutoff productivity level will produce in equilibrium. Therefore the equilibrium productivity distribution of plants choosing to operate, $\mathbf{m}(\mathbf{f})$, will be a truncation of $g(\mathbf{f})$:

$$\mathbf{m}(\mathbf{f}) = \begin{cases} \frac{g(\mathbf{f})}{1 - G(\mathbf{f}^*)} & \text{if } \mathbf{f} \geq \mathbf{f}^* \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

where \mathbf{f}^* is the productivity level such that $\pi(\mathbf{f}^*) = 0$, and $G(\mathbf{f}^*)$ is the cumulative distribution function of $g(\mathbf{f})$ evaluated at the cutoff productivity value. We can therefore rewrite the earlier expression for average productivity as

$$\tilde{\mathbf{f}} = \left[\frac{1}{1 - G(\mathbf{f}^*)} \int_{\mathbf{f}^*}^{\infty} \mathbf{f}^{s-1} g(\mathbf{f}) d\mathbf{f} \right]^{\frac{1}{s-1}} \quad (13)$$

Notice that given $g(\mathbf{f})$, the average productivity level depends only upon the cutoff.

The equilibrium value of \mathbf{f}^* is determined by two conditions that must hold in the steady state.

The first, as alluded to above, is that the plant with cutoff productivity \mathbf{f}^* must make zero profits. The second is a free entry condition, requiring that the expected payoff from entry be zero in equilibrium. Both of these conditions can be used to derive expressions of average plant profits as a function of the cutoff level; the intersection of these two functions determines \mathbf{f}^* in equilibrium. The expressions (6) for plant profits, (7) for revenue distribution, and (11) for average revenue and profits can be used to derive an expression relating average profit and cutoff productivity \mathbf{f}^* :

$$\bar{p} = \frac{\Pi}{M} = p(\tilde{\mathbf{f}}) = \frac{r(\tilde{\mathbf{f}})}{s} - f = \left[\frac{\tilde{\mathbf{f}}}{\mathbf{f}^*} \right]^{s-1} \frac{r(\mathbf{f}^*)}{s} - f \quad (14)$$

The equilibrium condition that the marginal plant earn zero profit implies that $r(\mathbf{f}^*) = s\mathbf{f}$. Thus we have the equilibrium expression relating average profits and the cutoff productivity level for the first condition:

$$\bar{p} = \left[\left(\frac{\tilde{\mathbf{f}}(\mathbf{f}^*)}{\mathbf{f}^*} \right)^{s-1} - 1 \right] f \quad (15)$$

I explicitly write average productivity as a function of the cutoff level to stress that \mathbf{f}^* enters into this function in two places.

The second condition requires that the expected value of entry be equal to zero. Recall that the entering firm, upon learning its productivity level, must decide whether to produce or exit before producing. Because productivity levels do not change over time, if it produces the establishment will make the same profit every period until it receives a killer shock (which has a probability of \mathbf{d} each period). Thus after the sunk cost, the value of entry is either zero if $\mathbf{f} < \mathbf{f}^*$, or (assuming no discounting other than knowledge of the possibility of an exogenous shock):

$$v(\mathbf{f}) = \sum_{t=0}^{\infty} (1-\mathbf{d})^t p(\mathbf{f}) = \frac{p(\mathbf{f})}{\mathbf{d}} \quad (16)$$

The expected value of entry is then the product of the probability and average value of successful entry minus the sunk entry cost.

$$v_e = (1 - G(\mathbf{f}^*)) \frac{\bar{p}}{\mathbf{d}} - f_e \quad (17)$$

Given that this must be zero in equilibrium, the second expression relating average profit and \mathbf{f}^* is obtained:

$$\bar{p} = \frac{df_e}{1 - G(f^*)} \quad (18)$$

This equation shows that the free entry condition requires average profits be increasing in f^* . This makes sense; as the cutoff productivity rises, the probability of successful entry decreases, requiring that the expected profit from successful entry increases to compensate for this fact.

The impact of changing f^* on the other equilibrium condition (15) is not so immediately apparent. Recall, however, that the first term in the bracketed expression is equal to the revenue ratio between the average productivity plant and the marginal plant. This climbs toward infinity as the cutoff productivity level goes to zero, because average productivity must always be positive. The behavior of this ratio as the cutoff productivity level increases depends on the properties of $g(f)$, but for common distribution functions, it will monotonically decrease toward one as f^* rises. Thus the requirement that the marginal plant make zero profits implies a negative relationship between average profits and the cutoff productivity level. This, too, is as expected. Since the marginally efficient plant earns zero profits, lower average profit levels imply that the productivity level required to operate in the black must increase. One can demonstrate that the properties of these two functions ensure a unique equilibrium cutoff productivity level f^* (see Melitz (1999)).

Notice that the elasticity of substitution s affects this level only through one function: that which is derived from the requirement of zero profits at the cutoff productivity level (15). An increase in s shifts this function up (average profit increases for all possible cutoff productivity levels). Because the free-entry condition implies average profits are an increasing function of the cutoff productivity, an increase in substitutability raises both the cutoff productivity and the average profit level in equilibrium. This in turn further truncates the ex-ante productivity distribution, decreasing productivity variation across producers and raising the average productivity level. This is depicted in Figure 2.1. If factors such as transport costs, physical product differentiation, and brand-creating expenditures influence the value of r for a given industry's output (and through it, the value of s), the truncation point and hence the moments of the industry productivity distribution will vary with output substitutability. Those industries with factors contributing to less output market segmentation (lower transport costs, more product homogeneity, and fewer brand identifiers) should have productivity distributions with less dispersion and higher central tendency. This, of course, is exactly the process outlined intuitively in the

introduction, and is the implication I test empirically below.

The model shows that factors other than output substitutability also change the shape of the equilibrium productivity distribution. Increases in the probability of exit, d , shift the free entry curve up, increasing the plant-level productivity dispersion and decreasing average productivity. Intuitively, free entry requires that industries with lower survivability rates require higher profits (for a plant of given productivity) to compensate for this additional risk. This implies that lower-productivity plants which would be unprofitable in less risky industries will be able to stay in business, thereby broadening the productivity distribution's range and lowering its average. Increased entry costs also shift the free-entry curve up, increasing dispersion and lowering the central tendency of the productivity distribution. In a process analogous to that driven by high exit probabilities, the preservation of free entry in high entry cost industries requires higher profits for each given productivity level, allowing less productive plants to operate. On the other hand, increased fixed *production* costs shift the "zero-cutoff-profit" curve up, decreasing dispersion and raising the average. The intuition here is that plants that are marginally profitable at lower fixed production costs cannot operate as these costs rise.

A further influence on the plant-level productivity distribution, as Melitz (1999) proceeds to show (and in fact is the focus of his analysis), is the degree of trade exposure in an industry. Two implications he examines are the move from autarky to a trading environment, and a reduction in fixed trade costs given that trade is already occurring. A move from autarky to trade drives low productivity domestic plants out of business (because of the entry of high-productivity foreign competitors) and increases the market shares of high productivity domestic operators. These effects combine to result in a narrowing of the productivity distribution and an increase in its average. Reductions in existing fixed trade costs also tighten the distribution, since the influx of foreign competitors who can now enter forces less productive domestics to exit. The impact on average productivity, however, is ambiguous. This is because two opposing influences may be at work as trade costs are lowered. The exit of low productivity domestic plants always tends to raise the average level, but the entry into foreign markets (and the resulting larger market shares) of formerly domestic-market-only plants from the middle of the productivity distribution can move the average up or down, depending on the previous level of trade costs. The reader is referred to Melitz (1999) for a much more in-depth discussion of these issues.

II. Empirical Method and Data

Specification

I test the implications of the model by seeing if measurable differences in industry output substitution elasticities change moments of their respective plant-level productivity distributions, controlling for other influences on those moments. This implies a general empirical specification of the form:

$$y_I = \mathbf{b}_0 + X_{sl} B_s + X_{cl} B_c + \mathbf{e}_I$$

This specification requires that the plant-level productivity distribution moment y_I (either a measure of dispersion or central tendency here) for industry I is a function of a constant, a vector of variables X_{sl} which change the elasticity of substitution for the industry's output, a vector X_{cl} of other influences on the moments, and an industry-specific error term. I will discuss the specific components of each of these vectors below.

The productivity distribution moments in this study are computed for each four-digit industry from plant-level data from the 1977 Census of Manufactures. For dispersion measures, I use various interquantile productivity differences, scaled by the median productivity in the industry. I compute four dispersion measures to check the robustness of results. Two are measured in terms of labor productivity: establishment gross output divided by number of employees, and output divided by total hours. I construct output by adjusting the plant's reported value of shipments for inventory changes over the course of the year. Plant hours are reported production worker hours plus a value for nonproduction worker hours computed according to the method of Davis and Haltiwanger (1991), where the plant's number of nonproduction workers is multiplied by the average annual hours for nonproduction workers in the corresponding two-digit industry (calculated from Current Population Survey data). In addition to these labor productivity measures, two total factor productivity values are computed for each establishment. Both follow the typical form:

$$tfp_i = y_i - \mathbf{a}_l l_i - \mathbf{a}_k k_i - \mathbf{a}_m m_i$$

where the lower-case letters indicate logarithms of the establishment-level TFP, output, number of

employees, capital stock, and materials inputs (energy inputs are included in materials). The two TFP measures differ by the manner in which the factor elasticities a_j are computed. One measure uses an average of 1972 and 1977 industry *revenue* shares, while the other utilizes the mean industry *cost* shares over the same two years. Expenditures are reported in the CM for labor and materials purchases; capital expenditures are computed for the cost-share basis by multiplying reported plant equipment and building stocks by their respective capital rental rates for each plant's corresponding two-digit industry; these rental rates are from the Bureau of Labor Statistics. While the CM contains data for every manufacturing establishment in the nation, approximately one-third of the plants are Administrative Record (AR) cases, which means all input data except the number of employees and total payroll are imputed. Therefore I use only non-AR plants in 1972 and 1977 to compute average industry shares.

Dispersion measures are scaled by the median productivity level in the industry to prevent pure scale differences between industries—mainly a factor in labor productivity measures due to capital intensity variation differences across industries—from causing productivity variation that is neither within the confines of the model nor very relevant to the hypothesis in this chapter. Using several interquantile differences as dependent variables allows robustness checks, as well as an ability to discern if output market segmentation affects segments of the plant-level productivity distribution asymmetrically. I use ordinal distribution moments rather than more common candidates, such as the coefficient of variation, because productivity values computed with plant-level data are especially vulnerable to influence of outlier plants. These are often a result of measurement error or miscoding in the original data. The few specifications I tried using labor productivity coefficients of variation did indeed yield erratic and inconclusive results.

For regressions with central tendency measures as the dependent variable, I use the median total factor productivity in the industry (both the revenue- and cost-share versions). I exclude labor productivity measures here because wide inter-industry variation in capital intensities causes average level differences that are outside the model. Unlike the dispersion measures, it is not possible to remove scale effects. TFP is much less susceptible to such problems. Median TFP levels are used rather than averages to counteract outlier effects.

Output Substitutability Factors

I would ideally like to regress these moments on values of the industry output substitutability s . However, industry output substitution elasticities are obviously not measurable, and it would be nearly impossible to estimate values given the sheer number of industries and products produced. My strategy is to instead fill X_{st} with measurable variables that influence output market segmentation. This allows indirect sizing of substitutability's impact on industry plant-level productivity distributions. Further, it will facilitate a partial decomposition of this effect into more fundamental components. For the intuition behind selection of variables for X_{st} , I return to the discussion above on the dimensions of market segmentation.

To capture between-industry differences in physical product variety, I take advantage of a CM variable that indicates the fraction of plant output composed of the plant's primary product (according to the seven-digit SIC product classification system). Industries with a large amount of physical variety in their products (distinctions in product classifications are largely made on the basis of physical differences) are likely to have plants which produce more than one physical product type. Conversely, plants in an industry having one product category will manufacture only one product type. Thus a higher average primary product specialization ratio (PPSR)¹ which I compute by industry using all non-AR plants in the 1977 CM—implies less physical product differentiation, a higher value of s , and less output market segmentation. As the model indicates, this should result in less dispersion in the productivity distribution and a higher mean productivity level. We should therefore expect a negative coefficient on PPSR in the dispersion regressions and a positive value in the level model.¹

Market segmentation in output markets can also be shaped by brand recognition and specialized sales efforts. In order to account for such influences, I utilize industry sales expenditures reports from the Federal Trade Commission's 1976 Line of Business report. This report has data for a large number of industries on selling expenditures as a percentage of total sales in that year. Reported

¹ The specialization ratio is not a perfect measure of physical product differentiation. It is possible that some industries with large product variety merely divide their output among a greater number of specialized establishments rather than producing several products at one plant. However, a low *PPSR* is certainly a sufficient condition for physical product differentiation, if not a necessary one.

sales expenditures numbers include not only media advertising, but also any other expenditures that are intended to influencing potential customers to purchase product, such as salaries for a sales and promotion staff. This ratio, SELLPCT, is a more complete measure of product differentiation efforts, because all such expenditures (including those beyond media advertising) are intended to differentiate one firm's products from another's. The FTC groups its data into roughly 200 industry groups. Given that there are about 450 four-digit SIC industries in 1977, many FTC industry categories consist of more than one SIC industry (the Line of Business report includes a concordance). Still, Table 1 indicates that there is considerable cross-sectional variation in selling expenditure shares. The expected sign for SELLPCT depends largely on one's view of the influence of advertising. The industrial organization literature is divided on this question.² One set of papers argues that advertising serves to confuse customers and create artificial product differentiation, largely along the lines of the branding motive discussed previously. This view holds that industries with higher ratios of media advertising and sales efforts expenditures to total sales should exhibit more product differentiation, thereby decreasing s , increasing long-run productivity dispersion, and lowering average industry productivity. This would be reflected by a positive coefficient in the dispersion model and a negative estimate in the levels equation. The opposing strand of literature contends that advertising serves to educate consumers about which products are superior. Selling expenditures under this view should allow the more productive firms to take market share away from less productive establishments, decreasing dispersion and increasing average productivity levels within the industry. If this effect dominated in practice, we should see a negative coefficient in the dispersion regressions and positive in levels. Of course, it is possible that both effects act simultaneously at roughly equal magnitudes, which will yield estimates showing no overall influence.

Output market segmentation can also arise geographically. Transport costs prevent plants from practicably selling their output beyond certain shipment distances. These distances, of course, depend on the magnitude of the transport costs. I use two methods to measure how large transport costs are for an industry; both use data from the 1977 Commodity Transport Survey (CTS). The CTS contains an enormous amount of information on the shipment of manufacturing products at a very detailed

² See the discussion in Tirole (1988) for specific citations.

product level (most data is presented at the five-digit product class level).³ Included in this survey are, for each five-digit product class, the average dollar value per pound of shipments and a decomposition of total tons shipped by distance category. I construct two indirect measures of transport costs using this information.

The first transport cost metric, VALUELB, simply aggregates the value-to-weight ratios (value densities) of all product classes in a given four-digit industry, weighted by the total tons shipped of each product class in the industry, and then takes the natural logarithm of the result.⁴ There is an obvious relationship between the value of shipments per pound and output transportability. Goods which are very valuable compared to their weight are much more economical to ship great distances, as the relative costs of transport decline with value density. Other goods—concrete blocks, for example—are so heavy relative to their value that it is not economically feasible for their production plants to ship product beyond a local area. Industries with high values of VALUELB should then have output markets that are less geographically segmented, and as such have more substitutability between the output of industry establishments. Low values imply greater spatial product differentiation. The coefficient on VALUELB should then be negative for the dispersion regressions and positive for the level models.

The second measure of geographic market segmentation utilizes CTS data on shipment sizes by distance. For each product, the CTS reports the number of tons shipped within each of seven distance-from-production-site categories.⁵ Product shipment data are combined to determine a shipment pattern

³ The CTS is incredibly detailed for public use data; as far as I am aware, nothing approaching its level of detail has been published (or possibly even collected) since. Because it has not been repeated, I am limited to using 1977 as my sample year; the CTS is, in effect, the binding data constraint in this study. Most of the plant and industry data needed for this study was collected in 1976 or 1977, so changes in industry patterns over time should have a minimal influence on the results.

⁴ While there is close correlation between the CTS product categories and the corresponding four-digit SIC industries that contain them, they do not perfectly match. Using published descriptions of product types found in each industry, I was able to aggregate products into their corresponding industry for nearly every SIC four-digit classification. Shipment data for the ordinance industries (SICS 3282-4,9) was not available; these industries were dropped from my sample. A concordance is available from the author upon request.

⁵ The categories are as follows: less than 100 miles, 100-199, 200-299, 300-499, 500-999, 1000-1499, and over 1500 miles.

for each four-digit industry. Industries which ship a greater percentage of their output shorter distances have output markets that are more geographically segmented. I use the shipment distance data to construct a measure of the geographic size of the industry output markets, LOCAL. This measure is a weighted sum of the percentage of output (by ton) shipped by distance categories. The weights are constructed by first computing the average shipment distance within each distance category, assuming shipments are distributed uniformly throughout the category. For example, the average distance for the 100-199 miles category is 150 miles, and it is 750 miles for the 500-999 miles division. (I use 1750 miles for the over-1500 miles shipments.) The sum of these seven lengths is divided by each category's average distance to get preliminary weights for the respective categories. Finally, these preliminary weights are normalized so that the weight of the less than 100 miles category is equal to one.⁶ We might expect that industries which ship a greater proportion of their output to areas near their plant of manufacture to be more geographically segmented and have less product substitutability across space.⁷ LOCAL is therefore expected to have a positive coefficient in the dispersion regressions and negative effect on the median productivity level in an industry.

Use of LOCAL as a measure of transport costs requires a caveat. If an industry's customers are geographically concentrated and industry plants choose to operate near their customers, it is possible that an industry could have high output substitutability despite a small average shipment radius. This is not a large issue for consumer goods industries, whose buyers are distributed throughout the country, but it may be for some industries which serve as suppliers of intermediate goods to specialized downstream buyers. To control for this possibility, I include a measure of industry geographic concentration GEOGCONC from Ellison and Glaeser (1997) in some of my specifications. Unfortunately, the Ellison-Glaeser industry data is available by the 1987 SIC system rather than the modified 1972 SIC system used in my sample year. To convert the original data to the earlier classification, I first match each plant to the concentration index value for the industry that it would be in

⁶ The weights are: less than 100 miles, 1; 100-199 miles, 0.333; 200-299, 0.2; 300-499, 0.133; 500-999, 0.067; 1000-1499, 0.04; and over 1500 miles, 0.029.

⁷ Continuing with the above example, plants in the concrete block and brick industry (SIC 3271) ships 94.2% of their output less than 100 miles.

under the 1987 system (this information is available in my data). I then take an average of these values over every 1972 industry to compute a concentration measure for industries in my sample. This method preserves a good portion of the relative concentrations for industries that were consolidated, split, or otherwise reclassified, while still allowing unchanged industries to have the same concentration index in 1977 as they would under the 1987 system.⁸ Given a value for LOCAL, we should expect higher values of GEOGCONC to lead to increased substitutability (because closely-spaced plants are able to ship to largely the same customer base, even in low-transportability industries). Hence the dispersion coefficient on GEOGCONC is expected to be negative, and its level estimate should be positive.

Other Influences on the Productivity Distribution

Of course, cross sectional differences in the ex ante productivity distribution $g(n)$ influence industries' productivity distributions. To the extent that these are reflected in output scale differences, I account for this influence by scaling dispersion moments to the median productivity level in the industry. I also test for robustness of the dispersion results by using various interquantile differences, which allows me to determine effects across different subsets of the productivity distribution. While clearly not a flawless solution, these steps should remove a sufficient amount of the influence of different ex-ante distributions across industries.

In an attempt to control for these other influences on industry productivity distributions, I use a number of proxy variables collected and constructed from an array of sources. It is not apparent beforehand whether failing to control for these other factors will bias the substitutability factor coefficients, as all are plausibly uncorrelated with market segmentation. However, adding proxies for these other effects allows further testing of the model's implications independent of any market segmentation consequences.

To control for the influence of exit probability on the productivity distribution, I construct a measure of exit rates for each industry. The model simplifies the exit process by assuming an exogenous

⁸ Of course, spatial patterns in industry certainly changed from my sample year to the 1992 sample date of the Ellison-Glaeser study. However, given the important influence of natural advantages (which do not vary in location over time) on manufacturing plant location found by Ellison and Glaeser, it is likely that any changes were reasonably small.

probability of exit that is not correlated with productivity. We may expect heterogeneous establishments to react differently in practice to their common industry exit rate. While neither the theoretical model nor the empirical model explicitly captures this process, it is very likely that industry exit rates are highly correlated with plants' perceived exit probabilities in the cross section. Therefore, a single measure of exit probability per industry should be able to account for this effect even if all plants within the industry do not respond to the measure equally. To construct an exit probability measure, I combine data from both the 1972 and 1977 Census of Manufactures, computing the fraction of industry plants present in 1972 that are also present in 1977. While this five-year survival rate, SURVRT, is not a forward-looking measure unless exit likelihoods are constant, it should be highly correlated with measures derived from more complicated expectations models. Increases in SURVRT are expected to decrease productivity heterogeneity, yielding a negative coefficient in the dispersion regressions and a positive estimate in the level equations.

Controlling for fixed production and entry costs can be a difficult matter empirically, because it is not clear which fixed costs faced by plants can be classified as entry costs and which as production costs. In the model, fixed entry costs are paid before the plant learns of its productivity level. If a plant learns its productivity level only after production begins, then any production-related overhead could be classified as either a fixed production cost or an entry cost (at least in the first year of production). In this framework, it would be difficult to discern the separate influences of the two cost types, as they move the productivity distribution in opposite directions. Thus using observables which could plausibly be linked to either cost structure could yield inconclusive results. If, on the other hand, a plant learns its productivity level before starting production, it is not immediately clear what would be an appropriate measure of entry costs; most data, after all, is collected after production is underway. I attempt to reconcile these factors by assuming that entry costs are related to post-production observables. Specifically, I follow the method employed by Sutton (1991), who constructs measures of sunk entry costs relative to market size for several food and drink industries. This measure, SUNKCOST, is equal to the proportion of industry's median-plant output to total industry output, multiplied by the ratio of total industry capital stock to industry output.⁹ While this measure benefits from its effectiveness in a

⁹ The ratio of median plant output to total industry output is often used in the empirical industrial

previous use, problems could still remain when interpreting its effect. The measure is closely related to the median plant-level capital stock in the industry, which may be a measure of fixed *production* costs if there is any overhead capital required in the production technology (this may be especially true of building stocks). Using a measure which is correlated with both entry and production fixed costs makes the expected sign of the coefficient indeterminate in the productivity dispersion and level regressions. Therefore appropriate care should be taken when interpreting the estimate of SUNKCOST.

To further address the fixed costs measurement problem, I use a measure of industry fixed production costs that is plausibly uncorrelated with entry costs. This value is FIXLAB, the average ratio of non-production workers to total employment at plants within the industry. Overhead labor, a production-specific fixed cost, is likely to be closely correlated with the share of non-production employees. Any comovement between this fixed production costs measure and entry costs should be only through pure scale effects, because it is unlikely that overhead labor is otherwise closely related to entry costs. Caution is still warranted in interpreting the FIXLAB coefficient, because the degree of influence that scale effects may have is unknown. It should be noted that including these fixed cost measures controls for their influence on output substitutability even though there may be difficulties in precisely decomposing their effects into entry and production components.

To measure an industry's trade exposure and cross sectional differences in fixed trade costs, I use both import- and export-based metrics computed from the Feenstra (1997) trade data. Industry import penetration, IMPPEN, is defined as the value of imports of a given industry's product(s) divided by the sum of these imports and the value of domestic production in the industry. Export intensity EXPINT is computed as industry exports divided domestic output for the sector. Obviously, larger values of either variable should coincide with greater trade exposure and lower fixed trading costs. Therefore we should expect these variables to have negative influences on productivity dispersion. As mentioned above, however, their expected effect on the central tendencies of industry productivity distributions cannot be signed.

organization literature as a measure of minimum efficient scale.

III. Results

I begin by regressing the plant-level productivity distribution moments on each of the output market segmentation factors separately. The results are presented in Table 2. Panel A shows the coefficients obtained by regressing each of the four productivity dispersion measures on the respective substitutability variables. (Although the coefficients are listed in columns under the dispersion measures, the factor coefficients are for single-variable [besides a constant] regressions.) The measure of plant-level dispersion for each productivity measure is the log of the ratio of the industry's interquartile productivity range to its median. Thus coefficients can be interpreted as percentage changes in productivity variability caused by the explanatory variables.

The results are consistent with my hypothesis; factors that plausibly increase the between-plant substitutability of an industry's output decrease the dispersion in the plant-level productivity distribution. Increases in VALUELB (the average value per pound of an industry's output) and PPSR (the average fraction of plant output accounted for by the plant's primary product), which should be positively correlated with substitutability, coincide with decreases in productivity dispersion. LOCAL (average shipment distance) has the expected positive correlation with dispersion. These findings are consistent across all four measures of market segmentation and the four productivity measures. In the 12 cases with expectations of estimates' signs (i.e., all but the regressions with SELLPCT—the industry ratio of advertising/sales-force expenditures to output), all are signed as predicted. Further, ten of these are significant at the 5% level, and eight at the 1% level. These single-factor exercises indicate that productivity dispersion and output market segmentation are undoubtedly correlated, and in directions consistent with the theory.

Regressing the median TFP level in each industry on individual output market segmentation factors yields largely the same consistent and intuitive results found in the dispersion regressions. This can be seen in panel B of Table 2. Value density and the average product specialization ratio in an industry are positively related to the median profit level, while LOCAL is negatively correlated. The coefficients from these regressions are signed as expected for both productivity measures, and are significant at the 1% level in four of six cases (one of the other two estimates is significant at the 10% level). The only contradictory results seen in these regressions are those using SELLPCT: both

coefficients are signed positively. While this does not contradict theory, which predicts an ambiguous effect given the possible dual function of sales expenditures discussed above, it is inconsistent with the results from the dispersion regressions. If sales expenditures serve to increase product differentiation, as it appears to do in the dispersion regressions, they ought to decrease median productivity rather than raise it.

I next regress the moments on all substitutability variables simultaneously. Because VALUEL and LOCAL both characterize geographic segmentation, I do not include them together in regressions. The outcome of these regressions is presented in Table 3. Again the results are striking in their consistency. Looking at the dispersion regression coefficients in Panel A shows that the estimates for VALUEL, LOCAL, and PPSR all have the expected sign, and are nearly always significantly estimated at standard confidence levels. These are also economically relevant magnitudes: a doubling of value density (industry output values range from \$0.01 to \$150 per pound) corresponds to a decline in the dispersion of labor productivity by roughly 2%, and by 3.5% for TFP. Industries shipping all output less than 100 miles (i.e., a value of LOCAL equal to one) have about 20% more labor productivity dispersion and 14% greater TFP dispersion than industries with average shipping practices, all else being equal. A completely specialized industry (PPSR=1) has about 7% less variability in labor and total factor productivity than an industry of average product specialization (PPSR=0.896). GEOGCONC, which was included to control for the influence of downstream-sector geographic concentration on LOCAL, has positive coefficients. This is counterintuitive; given an average shipment radius, industries with production facilities located closely together should have higher substitutability in output, and thus less dispersion. As in the single-variable cases, SELLPCT is found to have a significant and positive influence on dispersion; each one-percentage point hike in the fraction of total revenue devoted to sales expenditures increases labor productivity dispersion by roughly 0.5% and TFP dispersion by 1%.

The observed factors on substitutability explain anywhere from 3 to 7% of across-industry labor productivity dispersion differences, and from 8 to 13% of movements in plant-level total factor productivity variation from industry to industry. Clearly, these factors alone cannot explain the full extent of productivity heterogeneity in an industry. This is not surprising, given all of the variation in

technological (supply-side) influences across industries that shape productivity distributions. Output substitutability is surely an economically relevant part of the story, however. Furthermore, as I argued in the introduction, there are also non-measurable market segmentation influences. The results presented here imply a likelihood that these other factors affect industries' plant-level productivity distributions in the same manner as their measurable counterparts. The combined effect could be quite large.

As seen in Panel B of Table 3, the combination of all substitutability factors also shows the expected influence on median productivity levels across industries. VALUELB is positive and highly significant in both of its regressions. A doubling of value density corresponds with roughly a 10% increase in the TFP of an industry's median plant. PPSR, too, is positively and significantly related to median productivity. Its estimates imply that a perfectly specialized industry will be about 10% more productive than one of average specialization ratio. LOCAL is correctly signed but less precisely estimated, indicating that industries with completely locally-operating plants have a median productivity that is 15-30% lower than the average industry. GEOGCONC continues to offer counterintuitive coefficients that I am at a loss to explain at this point. SELLPCT here again contradicts its measured influence in the dispersion regressions with three of four estimates being positive. However, these estimates are less precisely estimated than in the dispersion models. The combination of the output market segmentation factors explains anywhere from 3.3 to 26.7% of the variance in median productivity levels depending on the productivity measure and explanatory variables.

I present estimates from the complete model in Table 4. Panel A shows the dispersion regression results. It is apparent that the hypothesized influence of market segmentation on heterogeneity found above was not a spurious result caused by correlation of the substitutability factors with other determinants of productivity dispersion. VALUELB, LOCAL, and PPSR retain their expected signs and are significantly estimated in every case. SELLPCT still has a positive influence on dispersion, but only significantly so for TFP measures. The magnitudes of these estimates are not drastically or systematically different from those shown in Table 3.

The other influences on heterogeneity indicated by the model have much more ambiguous and less statistically significant coefficients. The survival rate SURVRT appears to have no discernable effect on heterogeneity; its coefficients change from positive to negative depending on the productivity

measure and are never significantly different from zero. It is unclear whether this implies that SURVRT is a poor measure of what the model's exit rate d , or if the model's specification of an exogenous exit probability unrelated to productivity is too far from reality (this is quite likely). SUNKCOST, the measure of entry costs, is significantly estimated in the majority of regressions, but its coefficient also changes sign depending on whether labor or total factor productivity is used as the dependent variable. (Recall that fixed entry costs should be positively related to productivity dispersion.) FIXLAB, the proxy for fixed production costs, is erratically signed and never estimated to be different from zero. As discussed above, the fact that the two fixed costs variables may be correlated with both entry and production fixed costs, and that these two cost types have opposing affects in the model, may be responsible for the inconsistent results. Trade exposure seems to have an influence on heterogeneity only through import intensity. The coefficients on EXPINT are positive in seven of eight cases, though always imprecisely estimated. The industry import penetration ratio IMPPEN has a positive influence on dispersion and is significantly estimated in every case but one. The trade coefficients imply that trade-intensive industries have a broader productivity dispersion than those dominated by domestic plants. This seems to be counterintuitive at first glance. Perhaps imports serve to decrease substitutability by introducing more product differentiation and thereby increasing dispersion. Alternatively, there is a possibility that the IMPPEN results are being driven by reverse causation: foreign competitors may be more likely to enter industries that already exhibit a large amount of productivity dispersion, as their ex-ante probability of successful entry should be higher in such industries, *ceteris paribus*.

The explanatory power of the full dispersion model ranges from roughly 7 to 18%, depending on the included variables. Consistent with earlier patterns, the model tends to have better fit when TFP measures are used rather than labor productivity values.

The results of the full productivity level model shown in Panel B of Table 4 echo the findings of the dispersion regressions for those variables influencing s . A strong explanatory role is found for value density and product specialization, and to a lesser extent, geographic market area. The magnitudes of the estimates are also similar to those in Table 3. And as before, the role of sales expenditures in determining industry productivity levels is murkier.

For the other factors, the estimates for the productivity level equations are a bit more consistent with my priors. While there seems again to be no influence of exit probabilities on industry productivity distributions, the fixed cost variables offer coefficients of the expected signs and are for the large part significantly estimated. Fixed entry costs are negatively related to median productivity levels as predicted by the theory, and fixed production costs are found to have the opposite effect. Trade exposure variables still yield little information; the signs of both IMPPEN and EXPINT are quite erratic and statistically insignificant in six of eight cases.

The explanatory power of the model is considerably higher when the cost-based TFP measure is used as the dependent variable. This is consistent with the notion that revenue-based measures, which calculate capital's share as a simple residual, are more noisy. The cost-based values have less spurious variation and as such better conform to the predictions of the model.

Finally, I check the robustness of the full model's results to the productivity dispersion measure. The outcomes are presented in Table 4a. Panel A repeats the dispersion regressions using the log of the 90-10 interquantile range, and Panel B uses the log of the 95-5 percentile range. The estimated influences of the market segmentation factors on productivity dispersion are consistent with those discussed earlier. The productivity span used to measure dispersion does not seem to significantly influence the results.

A note of caution is warranted regarding these empirical results. All empirical testing was done using moments of *measured* productivity distributions. The theory conceptualizes productivity as efficiency in production; i.e., influencing input costs incurred in producing a unit of output. Empirical productivity measures, however, are not so cleanly obtained. Because plant output is measured in terms of revenue (due to a lack of comprehensive physical output data or plant-specific deflators), any plant-level price variation within an industry enters into plants' output measures. Thus price dispersion will create variation in productivity measures that is in addition to any differences in production efficiency.

This means, for instance, that industries with a large amount of physical product differentiation may have measured productivity dispersion that is greater than the true differences in technological efficiency across its plants. This is the case because the across-plant variation in physical characteristics of output will likely be embodied in price differences between plants. The increase in measured

productivity dispersion seen in industries having greater physical product differentiation may then be an artifact of price dispersion rather than technological dispersion. Separating the effects of efficiency from price differences is unfortunately not possible without detailed output product data or plant-specific deflators.

However, within-industry price variation should affect only the dispersion moments of the industries' plant-level measured productivity distributions. It should not influence between-industry differences in average productivity levels. Given that I found measurable substitutability factors to have the expected correlation with movements in the central tendency of industry productivity levels, it is likely that output market segmentation does indeed shape the distribution of productive efficiencies within and across industries.

IV. Conclusion

On whole, the results are consistent with the hypothesis that the degree of output market segmentation in an industry is related to moments of its plant-level productivity distribution. Factors plausibly correlated with high substitutability are shown to have a negative correlation with productivity dispersion and a positive influence on median productivity. This is shown to be the case both including and not including controls for other influences on the industry's productivity distribution.

Further, these substitutability/segmentation factors explain a nontrivial fraction of the total interindustry variation in productivity moments. This suggests that, while technological factors doubtlessly play a role in causing productivity dispersion, demand-side influences are also important. Exploring the specific output market mechanisms driving these results may be a fruitful path for further research.

References

- Davis, Steven J., and Haltiwanger, John. "Wage Dispersion between and within U.S. Manufacturing Plants, 1963-1986." *Brookings Papers on Economic Activity, Microeconomics 1991*, 115-80.
- Dixit, Avinash K., and Stiglitz, Joseph E. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, 67(3), 1977, pp. 297-308.
- Dunne, Timothy, Roberts, Mark J., and Samuelson, Larry. "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4), 1989, pp. 671-98.
- Ellison, Glenn, and Glaeser, Edward L. "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach." *Journal of Political Economy*, 105(5), 889-927.
- Feenstra, Robert C. "NBER Trade Database, Disk 3: U.S. Exports, 1972-1994, with State Exports and Other U.S. Data" NBER Working Paper no. 5990, April 1997.
- Hotelling, H. "Stability in Competition." *Economic Journal*, 37(1), 1929, pp. 41-57.
- Melitz, Marc J. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." Mimeo, University of Michigan, November 1999.
- Salop, Steven C. "Monopolistic Competition with Outside Goods." *Bell Journal of Economics*, 10(1), 1979, pp. 141-56.
- Sutton, John. *Sunk Costs and Market Structure*. Cambridge, MA: MIT Press, 1991.
- Tirole, Jean. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press, 1988.
- U.S. Bureau of the Census. *Commodity Transport Survey*. Washington, DC. 1977.
- U.S. Office of Management and Budget. *Standard Industrial Classification Manual*. Washington, DC. 1972.
- Weitzman, Martin L. "Monopolistic Competition with Endogenous Specialization." *Review of Economic Studies*, 61(1), 1994, pp. 45-56.

Table 1: Summary Statistics of Regression Variables

Variable	Mean	Std. Dev.
VALUELB	0.352	1.641
LOCAL	0.412	0.163
GEOGCONC	0.054	0.077
PPSR	0.896	0.068
SELLPCT	0.092	0.059
SURVRT	0.715	0.154
FIXLAB	0.232	0.093
SUNKCOST	2.34×10^{-3}	9.54×10^{-3}
IMPPEN	0.076	0.104
EXPINT	0.059	0.069

N=443

Table 2: Regression Results—Univariates on Market Segmentation Factors**A. Dispersion Regressions (log[Interquartile Range ÷ Median Productivity] as Dependent Variable)**

Factor	Productivity Measure in Dependent Variable			
	LP (Employees)	LP (Hours)	TFP (Revenue)	TFP (Cost)
VALUELB	-0.028** (0.010)	-0.021* (0.010)	-0.043** (0.010)	-0.042** (0.009)
LOCAL	0.325** (0.102)	0.272** (0.099)	0.169 (0.118)	0.186* (0.081)
PPSR	-0.608** (0.220)	-0.645** (0.215)	-0.390 (0.275)	-0.657** (0.222)
SELLPCT	0.420 (0.274)	0.326 (0.273)	1.035** (0.326)	0.589** (0.224)

B. Level Regressions (Median Productivity as Dependent Variable)

Factor	Productivity Measure	
	TFP (Revenue)	TFP (Cost)
VALUELB	0.147** (0.011)	0.139** (0.013)
LOCAL	-0.538 (0.379)	-0.237 (0.131)
PPSR	1.055** (0.387)	1.103** (0.400)
SELLPCT	1.709* (0.789)	0.914* (0.379)

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.

Table 3: Regression Results—Substitutability Factors Only

A. Dispersion Regressions (log[Interquartile Range ÷ Median Productivity] as Dependent Variable)

Factor	Productivity Measure in Dependent Variable							
	Labor (emp)		Labor (hours)		TFP (revenue)		TFP (cost)	
VALUELB	-0.032**		-0.023*		-0.055**		-0.048**	
	(0.011)		(0.010)		(0.009)		(0.008)	
LOCAL		0.385**		0.327**		0.225		0.254**
		(0.104)		(0.101)		(0.116)		(0.080)
GEOGCONC		0.489*		0.402		1.157**		0.119
		(0.214)		(0.212)		(0.313)		(0.265)
PPSR	-0.590**	-0.780**	-0.636**	-0.789**	-0.377	-0.637*	-0.623**	-0.779**
	(0.215)	(0.220)	(0.212)	(0.215)	(0.255)	(0.256)	(0.205)	(0.214)
SELLPCT	0.690*	0.585*	0.545*	0.482	1.435**	1.001**	0.969**	0.762**
	(0.278)	(0.286)	(0.277)	(0.282)	(0.310)	(0.342)	(0.209)	(0.227)
R ²	0.045	0.070	0.035	0.057	0.078	0.088	0.126	0.071

B. Level Regressions (Median Productivity as Dependent Variable)

Factor	Productivity Measure			
	TFP (revenue)		TFP (cost)	
VALUELB	0.138**		0.136**	
	(0.011)		(0.013)	
LOCAL		-0.486		-0.225
		(0.357)		(0.135)
GEOGCONC		-1.178*		-1.252**
		(0.464)		(0.281)
PPSR	0.669**	1.143**	0.787*	1.205**
	(0.300)	(0.386)	(0.339)	(0.372)
SELLPCT	0.738	1.581*	-0.058	0.903*
	(0.794)	(0.713)	(0.377)	(0.398)
R ²	0.072	0.033	0.267	0.090

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.

Table 4: Regression Results—Full Model

A. Dispersion Regressions (log[Interquartile Range ÷ Median Productivity] as Dependent Variable)

	Productivity Measure in Dependent Variable							
	Labor (emp)		Labor (hours)		TFP (revenue)		TFP (cost)	
VALUELB	-0.041**		-0.043**		-0.059**		-0.056**	
	(0.012)		(0.011)		(0.010)		(0.009)	
LOCAL		0.414**		0.377**		0.298*		0.335**
		(0.115)		(0.110)		(0.128)		(0.087)
GEOGCONC		0.491*		0.411*		0.964**		0.022
		(0.206)		(0.199)		(0.320)		(0.269)
PPSR	-0.645**	-0.805**	-0.607**	-0.749**	-0.615**	-0.782**	-0.566**	-0.729**
	(0.224)	(0.232)	(0.218)	(0.225)	(0.241)	(0.242)	(0.215)	(0.229)
SELLPCT	0.592*	0.481	0.422	0.343	1.697**	1.317**	1.056**	0.922**
	(0.289)	(0.301)	(0.281)	(0.291)	(0.323)	(0.379)	(0.212)	(0.239)
SURVRT	-0.061	-0.014	-0.079	-0.038	-0.024	0.031	0.034	0.066
	(0.108)	(0.103)	(0.106)	(0.102)	(0.106)	(0.105)	(0.085)	(0.085)
SUNKCOST	-3.814*	-2.535	-3.283	-2.172	5.402**	6.835**	3.025*	4.859**
	(1.794)	(1.726)	(1.796)	(1.686)	(1.554)	(1.782)	(1.414)	(1.754)
FIXLAB	0.177	0.186	0.341	0.344	-0.359	-0.257	0.182	0.125
	(0.180)	(0.192)	(0.179)	(0.189)	(0.192)	(0.217)	(0.151)	(0.176)
IMPPEN	0.375*	0.306	0.410*	0.356*	0.561**	0.395*	0.473**	0.381**
	(0.191)	(0.177)	(0.184)	(0.173)	(0.187)	(0.174)	(0.135)	(0.127)
EXPINT	0.166	0.119	0.309	0.286	0.274	-0.005	0.345	0.206
	(0.345)	(0.352)	(0.311)	(0.313)	(0.225)	(0.269)	(0.255)	(0.288)
R ²	0.067	0.083	0.067	0.083	0.134	0.133	0.187	0.130

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.

Table 4: Regression Results—Full Model (Continued)**B. Level Regressions (Median Productivity as Dependent Variable)**

	Productivity Measure			
	TFP (revenue)		TFP (cost)	
VALUELB	0.141**		0.139**	
	(0.013)		(0.014)	
LOCAL		-0.545		-0.230
		(0.372)		(0.145)
GEOGCONC		-1.060*		-0.965**
		(0.506)		(0.216)
PPSR	0.642*	0.995**	1.406**	1.685**
	(0.310)	(0.345)	(0.334)	(0.380)
SELLPCT	0.209	0.938	-0.548	0.321
	(0.744)	(0.657)	(0.333)	(0.398)
SURVRT	-0.430	-0.518	-0.016	-0.081
	(0.264)	(0.308)	(0.136)	(0.159)
SUNKCOST	-5.189*	-9.079**	-2.996**	-6.576*
	(2.066)	(3.165)	(0.329)	(2.762)
FIXLAB	0.713*	0.648	1.293	1.194**
	(0.348)	(0.392)	(1.713)	(0.387)
IMPPEN	-0.118	0.231	-0.424*	-0.033
	(0.231)	(0.220)	(0.187)	(0.183)
EXPINT	-0.755*	-0.107	-0.364	0.456
	(0.324)	(0.447)	(0.321)	(0.422)
R ²	0.085	0.050	0.348	0.174

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.

Table 4a: Additional Dispersion Regression Results—Full ModelA. $\log[90\text{-}10 \text{ Interquantile Range} \div \text{Median Productivity}]$ as Dependent Variable

	Productivity Measure in Dependent Variable							
	Labor (emp)		Labor (hours)		TFP (revenue)		TFP (cost)	
VALUELB	-0.026*		-0.022		-0.058**		-0.043**	
	(0.013)		(0.013)		(0.010)		(0.009)	
LOCAL		0.416**		0.368**		0.296*		0.358**
		(0.114)		(0.110)		(0.116)		(0.088)
GEOGCONC		0.275		0.348		0.878**		0.294
		(0.224)		(0.206)		(0.308)		(0.227)
PPSR	-0.477*	-0.611**	-0.511*	-0.629**	-0.608**	-0.771**	-0.633**	-0.782**
	(0.229)	(0.234)	(0.230)	(0.233)	(0.214)	(0.226)	(0.236)	(0.242)
SELLPCT	0.590	0.617	0.611*	0.617	1.654**	1.300**	0.980**	0.871**
	(0.317)	(0.328)	(0.311)	(0.318)	(0.280)	(0.338)	(0.063)	(0.241)
SURVRT	-0.202	-0.164	-0.259*	-0.223*	-0.183	-0.131	-0.234**	-0.196*
	(0.112)	(0.108)	(0.113)	(0.109)	(0.103)	(0.104)	(0.071)	(0.088)
SUNKCOST	-0.826	0.123	0.465	1.229	1.662	3.095**	0.378	1.748
	(1.653)	(1.438)	(1.654)	(1.519)	(1.001)	(1.147)	(1.331)	(1.189)
FIXLAB	0.111	0.086	0.175	0.171	-0.317	-0.229	-0.123	-0.138
	(0.199)	(0.207)	(0.193)	(0.203)	(0.170)	(0.199)	(0.153)	(0.165)
IMPPEN	0.301	0.288	0.193	0.177	0.575**	0.417*	0.509**	0.438**
	(0.195)	(0.184)	(0.194)	(0.183)	(0.217)	(0.199)	(0.169)	(0.156)
EXPINT	0.061	0.136	0.240	0.301	0.249	-0.014	0.225	0.153
	(0.329)	(0.302)	(0.302)	(0.277)	(0.215)	(0.263)	(0.259)	(0.268)
R ²	0.034	0.058	0.041	0.063	0.140	0.135	0.141	0.135

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.

Table 4a: Additional Dispersion Regression Results—Full ModelB. $\log[95\text{-}5 \text{ Interquantile Range} \div \text{Median Productivity}]$ as Dependent Variable

	Productivity Measure in Dependent Variable							
	Labor (emp)		Labor (hours)		TFP (revenue)		TFP (cost)	
VALUELB	-0.023		-0.017		-0.069**		-0.051**	
	(0.014)		(0.014)		(0.010)		(0.009)	
LOCAL		0.466**		0.445**		0.287*		0.373**
		(0.131)		(0.128)		(0.114)		(0.088)
GEOGCONC		0.354		0.374		0.844**		0.502*
		(0.231)		(0.222)		(0.300)		(0.229)
PPSR	-0.538*	-0.679**	-0.464	-0.592*	-0.519*	-0.697**	-0.662**	-0.828**
	(0.250)	(0.247)	(0.254)	(0.246)	(0.219)	(0.243)	(0.248)	(0.258)
SELLPCT	0.883*	0.932*	0.683	0.751*	1.620**	1.207**	0.932**	0.742**
	(0.366)	(0.371)	(0.366)	(0.368)	(0.276)	(0.337)	(0.232)	(0.255)
SURVRT	-0.162	-0.119	-0.233*	-0.192	-0.138	-0.085	-0.291**	-0.245**
	(0.117)	(0.116)	(0.115)	(0.113)	(0.097)	(0.103)	(0.091)	(0.088)
SUNKCOST	-3.122*	-2.228	-2.176	-1.481	-0.462	1.279	-1.621	-0.126
	(1.548)	(1.426)	(1.570)	(1.483)	(1.104)	(1.296)	(1.014)	(0.983)
FIXLAB	-0.039	-0.058	0.076	0.065	-0.308	-0.225	-0.211	-0.195
	(0.219)	(0.217)	(0.217)	(0.216)	(0.169)	(0.206)	(0.156)	(0.177)
IMPPEN	0.447*	0.444*	0.438*	0.446*	0.645**	0.459**	0.504**	0.402*
	(0.216)	(0.207)	(0.208)	(0.197)	(0.185)	(0.170)	(0.181)	(0.165)
EXPINT	0.096	0.208	0.206	0.341	0.440*	0.107	0.304	0.172
	(0.438)	(0.407)	(0.424)	(0.392)	(0.208)	(0.263)	(0.218)	(0.231)
R ²	0.037	0.065	0.035	0.065	0.156	0.123	0.152	0.143

*Denotes significance at the 5% level.

** Denotes significance at the 1% level.

Standard errors are White-consistent.